

Characterizing the folding core of the cyclophilin A — cyclosporin A complex II: improving folding core predictions by including mobility

J. W. Heal^{*†}, S. A. Wells[‡], R. B. Freedman[§] and R. A. Römer[¶]

[†]MOAC Doctoral Training Centre and Institute for Advanced Study, University of Warwick, Coventry, UK. [‡]Department of Physics, University of Bath, Bath, UK. [§]School of Life Sciences, University of Warwick, Coventry, UK. [¶]Department of Physics and Centre for Scientific Computing, University of Warwick, Coventry, CV4 7AL, UK.

Abstract

Determining the folding core of a protein yields information about its folding process and dynamics. The experimental procedures for identifying the amino acids which make up the folding core include hydrogen-deuterium exchange and Φ -value analysis and can be expensive and time consuming. As such there is a desire to improve upon existing methods for determining protein folding cores theoretically. Here, we use a combined method of rigidity analysis alongside coarse-grained simulations of protein motion in order to improve folding core predictions for unbound CypA and for the CypA-CsA complex. We find that the most specific prediction of folding cores in CypA and CypA-CsA comes from the intersection of the results of static rigidity analysis, implemented in the FIRST software suite, and simulations of the propensity for flexible motion, using the FRODA tool.

Revision :, compiled July 17, 2014

*Correspondence: jack.heal@bristol.ac.uk

Address reprint requests to Jack Heal, University of Bristol, School of Chemistry, Cantock's Close, Bristol, BS8 1TS, UK.

INTRODUCTION

Predicting folding cores

In Ref. (1) we determined the experimental folding core for the complex between the multifunctional protein cyclophilin A (CypA) and the immunosuppressant drug cyclosporin A (CsA) using hydrogen-deuterium exchange NMR (HDX). We showed that the established theoretical method of rigidity analysis implemented using the FIRST software provides a reasonable prediction of the experimental folding core (2–4). In addition to this, we generated the FIRST folding core (FIR) for unbound CypA, which broadly agreed with the previously published data from HDX experiments. Whilst impressive, the FIRST method is not perfect and does not capture the changes in the HDX folding core observed upon ligand binding. Indeed, FIR is larger in the absence of the ligand, a result which contrasts with what is observed experimentally. That there exists some discrepancy between theory and experiment could be expected since FIRST is used to make inferences about protein flexibility based upon a *static* crystal input structure only. In addition to flexibility, HDX is

dependent upon the surface exposure of the exchanging protons as well as protein motion (5, 6). With this in mind, we now seek to improve upon the predictive power of the rigidity analysis method by incorporating information on the surface exposure and dynamics of the protein. We change the hydrogen bond (HB) and hydrophobic tether (HP) networks so that these interactions are absent for surface atoms. With this modification we enhance the flexibility of the protein surface so that the remaining rigid residues may better correlate with the experimental folding core. Building upon this, we model the prospective dynamics of the protein using coarse-grained simulations. With these techniques we are able to probe the propensity for large-amplitude motions which may only be accessed over long timescales, such as the timescales of our HDX experiments.

Coarse-grained simulations of protein motion

The study of protein motion is of fundamental importance in structural biology due to its intimate link with

protein function (7). The most prevalent method for modeling protein motion computationally is through molecular dynamics (MD) simulations (8), yielding detailed atomic trajectories (9). MD simulations are known to be computationally intensive; yet some of the most biologically relevant motions take place on the ms-s timescale which remains inaccessible to MD at present (10). Methods of coarse graining may be applied to reduce computational demand and to enable simulation of larger amplitude protein motion (11). The constrained geometric simulation software FRODA (framework rigidity optimized dynamic algorithm) is a module within FIRST which rapidly simulates possible protein motion (10). It has been used to investigate protein-protein docking problems (12) as well as substrate recognition and conformational changes (13). Large-amplitude motions in proteins with hundreds or thousands of residues can be rapidly explored using a desktop computer on a timescale of minutes (14) — an improvement on the order of 10^4 over MD in terms of computational speed for motions over comparable distances. With FRODA we direct motion along a small number of low-frequency normal mode vectors (15–19), the superposition of which well describes large-amplitude protein mobility (14, 20, 21). The resulting potential for FRODA to access the rarely sampled, high energy protein conformers to which HDX is sensitive (22) ensures that it is a natural choice of simulation method for improving our folding core predictions.

MATERIALS AND METHODS

In addition to the FIR folding core, we now introduce four other possible theoretical estimates for the folding core. These include information on the surface exposure and dynamics of the protein in addition to rigidity. The PDB structure 1CWA (23) was used for all simulations of protein rigidity and motion, before and after removal of the ligand CsA. Initially we focus upon the CypA-CsA complex, before applying our analysis to the unbound CypA and comparing our results.

A modified bond network for FIR_B

The presence and strength of HBs within the constraint networks were determined based upon donor-acceptor geometry using FIRST. Bonds weaker than an energy cutoff value E_{cut} are excluded from the network. During a rigidity dilution E_{cut} is systematically lowered and HBs are consequently removed.¹ In addition to the standard bond network defined in this way by FIRST, we will

¹We emphasize that E_{cut} should be thought of as an arbitrary parameter and not as directly related to an energy scale for motion — such as $k_B T$ — in the structure (14).

also consider a modified network based on the surface exposure of atoms with HBs and HPs. By demanding that both interacting atoms in a HB or HP constraint are buried within the protein, i.e. not exposed on the surface, we enhance the flexibility of the protein surface. The folding core generated in this way consists of residues which are both rigid and buried within the protein. HPs are indirect, entropy-driven interactions (24) thought to contribute significantly to protein folding (25). HPs have not always been included in FIRST simulations, and their effect on modeling the flexibility and mobility of proteins is still not fully understood. In early papers using FIRST, HPs were not included in the bond network (26, 27), whereas more recently it is usual practice to include HPs and maintain them throughout rigidity dilution (28, 29) or to increase their number as E_{cut} is lowered (30, 31). In FIRST, HPs are modelled as flexible constraints, restricting the separation distance between atoms involved in the interaction, but not the angle between them. In this way, the interacting atoms are permitted to slip relative to each other (2, 32). HPs between carbon or sulfur atoms are typically included if the distance between these atoms is less than the sum of their van der Waals radii, r_v , plus a distance cutoff, D_{HP} . For carbon and sulfur, $r_v = 1.7 \text{ \AA}$ and 1.8 \AA respectively, and D_{HP} is typically set to 0.25 \AA (2, 29, 32). These distances will then allow to define the burial distance of a given atom (see below).

As a first improvement to FIR, we therefore modified the bond network in FIRST, requiring that each HP and HB is formed between two interacting atoms, both of which are buried within the protein. We implemented this bond network and performed a rigidity dilution on the structure 1CWA. As in Ref. (1), the folding core is determined from the dilution plot. In this case, the folding core consists of residues which are rigid and buried, and we refer to this measure as FIR_B .

Burial distances

Let us now define in detail how we distinguish buried and exposed residues. For each atom in the protein, a sphere of radius $r_v + r_w$ is drawn with the atom in the centre, where r_v is the van der Waals radius of the atom and $r_w = 1.4 \text{ \AA}$ is that of a water molecule; for nitrogen, $r_v(N) = 1.5 \text{ \AA}$. The points forming the sphere are then individually checked for contact with the neighboring atoms. If any point on the sphere is not in contact with a neighbor then that is a potential solvent position and the atom is labeled as being exposed. If an atom is not exposed, it is buried and its burial distance, R , is the shortest distance to an exposed atom. An exposed atom has $R = 0 \text{ \AA}$. For a given conformer, we calculated R

for each nitrogen atom in the amide backbone and used this as the value of R for the residue to which it belongs.

Coarse-grained mobility simulation for FRO

The coarse-grained elastic network model implemented using the ELNEMO software (17) allows for a rapid and accurate calculation of the normal modes of motion (16, 33). In Ref. (34) we reported results for directed motion along these normal modes to generate new conformations satisfying the constraints of the FIRST bond network. Following this work, we bias motion along both (parallel and antiparallel) directions of the ten lowest frequency non-trivial normal mode vectors, $m_7 - m_{16}$ with directed and random step sizes both set to 0.01 Å. These low frequency modes well capture the large scale motion of a protein (14). A total of 2000 conformers were generated for each mode and direction. Details of the algorithm can be found elsewhere (10). In each of our FRODA simulations, we use $E_{\text{cut}} = -2.0$ kcal/mol to determine the HB network. These simulations have been repeated at five different E_{cut} values ($-0.5, -1.0, -1.5, -2.0$ and -3.0 kcal/mol), but the nature of the results is not particularly sensitive to this (data not shown, see (35)).

We carried out simulations of motion using FRODA and calculated burial distances for all residues in the final (2000th) conformer from each of the FRODA simulations biased along a low frequency normal mode vector. The residues were classed as being buried ($R > r_v(N) \equiv 1.5$ Å) or exposed ($R \leq 1.5$ Å) in each of the conformers. Those which remained buried in more than half of the final conformations formed the FRODA folding core, which we refer to as FRO.

Secondary structure estimate SeS and the FIR+FRO combination F+F

As well as FIR, FIR_B and FRO introduced above, we consider two further theoretical folding cores, SeS and F+F. These are formed of the secondary structural units of helices and sheets (SeS) and the intersection of results between FIR and FRO (F+F). We choose to include SeS because of its simplicity and lack of need for any sophisticated computational analysis. The estimate F+F was included in order to test whether there is a distinct improvement upon combining data from static (rigidity, FIR) and dynamic (mobility, FRO) measures. In addition, we checked the intersection between FIR_B and FRO. This is missing just one residue Val12 from F+F, but is otherwise identical. Since Val12 is part of the HDX folding core, the FIR_B+FIR folding core makes a slightly weaker prediction than F+F in the present case.

Quantitative measures for comparing folding cores

Let N_{Ex} and N_{Th} be the number of residues contained in an experimentally and theoretically determined folding core, respectively. Clearly, $N_{\text{Ex}} = N_{\text{Th}}$ is a necessary condition for agreement of experimentally and theoretically estimated folding cores. However, not only the number of residues, but of course the agreement of the specific set of residues in both experimental and theoretical folding cores is most important. In order to capture this, let us define \mathcal{T} as the number of residues correctly identified by a theoretical prediction of the experimental folding core, i.e. \mathcal{T} is the number of true positives. For a perfect agreement, we have $\mathcal{T} = N_{\text{Th}}$ while the expected \mathcal{T} attained randomly is $N_{\text{Th}}N_{\text{Ex}}/N$. Here N denotes the total number of amino acids in the protein (3). We can define the specificity, α , and sensitivity, γ , of a theoretical folding core prediction as

$$\alpha = \frac{\mathcal{T}}{N_{\text{Th}}} \quad (1)$$

and

$$\gamma = \frac{\mathcal{T}}{N_{\text{Ex}}}. \quad (2)$$

The specificity α measures the fraction of residues identified by the theoretical method which are also part of the experimental folding core. The sensitivity γ shows the proportion of residues in the experimental core which have been correctly predicted by the theoretical method. A perfect correspondence between theory and experiment, $\mathcal{T} = N_{\text{Th}} = N_{\text{Ex}}$, would yield $\alpha = \gamma = 1$ while for a completely wrong identification, $\mathcal{T} = 0$, we have $\alpha = \gamma = 0$.

Another, previously defined (3), quantitative measure is the so-called folding core identification enhancement factor, ϵ . This is the ratio of \mathcal{T} and the number of residues expected to be identified by a random selection,

$$\epsilon = \frac{\mathcal{T}N}{N_{\text{Th}}N_{\text{Ex}}}. \quad (3)$$

A theoretical method with random probability of success has $\epsilon = 1$, and when $\epsilon > 1$, the match is better than random. We shall also compare our α and γ measures to ϵ .

RESULTS AND DISCUSSION

Effects of the modified bond network and the inclusion of mobility

We measured the surface exposure of each amide nitrogen in the structure 1CWA, cp. Figure 1. For the small protein CypA, we find that the majority of the residues are somewhat exposed to the protein surface,




Figure 1: The burial distance R for each residue of CypA in the static structure of the CypA-CsA complex. The blue dashed line at $R = 1.5$ Å is used to separate exposed residues (below the line) from buried residues (above the line). Below the horizontal axis, we schematically show the extent of the FIR_B folding core along the residues via coloured regions. Those residues which are part of FIR_B are coloured red; other rigid residues are coloured otherwise. The thin horizontal black line denotes residues that are flexible.

with $R < r_v(N) = 1.5$ Å. Since we measure the burial of the amide nitrogen and the length of the N-H bond is 1.5 Å, if the H atom is exposed to the surface, then the amide nitrogen has $R \leq 1.5$ Å. In Figure 1 we also show a superposition of FIR_B along the protein backbone as taken from the rigidity dilution plot. We note that most of the residues that are highly buried within the protein structure ($R \gtrsim 2.0$ Å) correspond to regions of the protein that are part of the folding core. This shows that the majority of those residues that are buried in the CypA-CsA complex are in fact rigid and buried.

Let us now see how our improved FIR_B folding core estimate is different from the FIR folding core of Ref. (1). Figures 2 (a) and (b) show a cartoon representation of the residues identified as belonging to the FIR_B folding core superimposed on the 3D structure of the protein. We also indicate in Fig. 2 (c) the FIR_B folding core along the protein chain as in Figure 1. This is similar to the cartoon representation of the comparison between the HDX folding core and the FIR estimate given in Ref. (1). As expected, FIR_B is smaller than FIR, since there are fewer constraints present throughout the rigidity dilution. The amino acids 12–18 and 41–47 are part of FIR but not part of FIR_B . These residues are close to the protein surface on the opposite side of the protein to the ligand binding site as can be seen in Figure 2(b).




Figure 2: (a) PyMOL cartoon of 1CWA with FIR_B in blue and the ligand, CsA, in black. Those residues which are part of FIR but not FIR_B are colored red, with the other residues of CypA in green (the colour scheme is similar to the one used in Fig. 1). (b) The surface of the protein is illustrated as a mesh, with CsA represented as spheres. The image has been rotated from its orientation in (a) by 180° about the 45° diagonal. (c) A linear representation by residue number of the rigid cluster decomposition at $E_{\text{cut}} = -1.061$ kcal/mol shown as thick colored blocks (rigid) or thin black lines (flexible). The mutually rigid red blocks show the amino acids which constitute FIR_B .

Figure 3(a) shows a qualitative comparison between FIR_B and the HDX folding core, depicted on the 3D structure of the CypA-CsA complex. We find that for most of the residues in the protein, the theoretical and experimental folding cores are in agreement. However, residues 138 – 143, forming the α -helix at the forefront of the image, are part of the HDX folding core, but are not part of FIR_B . With reference to the representation of FIR_B along the protein chain, as shown in Figure 2(c), we see that these residues are indeed rigid at E_{cut} corresponding to FIR_B , but they are not part of the largest rigid cluster. In this case, loosening the definition of the FIRST folding core to include residues that are rigid independently of the largest rigid cluster would capture this slowly exchanging helix as part of the rigid folding core.

In Figure 3(b) we show the FRODA folding core compared qualitatively with the HDX folding core. Notably, residues 138 – 143 in the helix are now part of the FRO folding core (but not of FIR_B). We emphasise that the results of this mobility-based analysis depend indirectly on the results of the rigidity analysis; FIRST identifies a constraint network and FRODA then explores the motion that is possible within those constraints. Therefore it is possible for a residue to be (i) rigid but exposed — thus lying in FIR or FIR_B but not in FRO — or (ii)




Figure 3: Theoretical folding cores (a) FIR_B and (b) FRO for the CypA-CsA complex compared with the HDX folding core. In both cases, residues are colored purple (part of both folding cores), blue (HDX folding core only), red (theoretical folding core only) or green (not part of either folding core). CsA is colored black.

non-rigid but well protected by burial within the protein — thus lying in FRO but not in FIR or FIR_B . This suggests that in order to predict the results of the HDX folding core, which depend upon surface exposure as well as dynamics (inherited from flexibility) a combination of the approaches employed by FIRST and by FRODA may be necessary.

In Figure 4, each of the five theoretical folding cores are compared with HDX folding cores for the CypA-CsA complex. In each case, folding core residues along the protein backbone are shown as bold blocks of colour. FIR , FIR_B and FRO have been discussed above, so we here briefly discuss the remaining folding cores definition SeS and F+F . We find that SeS consists almost entirely of slowly exchanging amino acids, although there are also many slowly exchanging residues which are not part of a secondary structure unit. The residues in FRO are also for the most part slowly exchanging. We note that some parts of the HDX folding core, for example residues 86–90, are poorly predicted by all of the methods. This region corresponds to a flexible, surface-exposed, unstructured region in the crystal structure, indicated by an arrow in Figure 5(a). Our methods, based around predicting rigid, buried and immobile regions, fall short of identifying this area as slowly exchanging. Figure 5(b) presents a graphical comparison between the HDX folding core and F+F on the structure of CypA. We observe that in the absence of the ligand, residues 86–90 are mostly absent from

both theoretical and experimental folding cores. The relatively small number of false positives (red) to false negatives (blue) demonstrate the high ratio of specificity to sensitivity for this method.

Quantitative analysis of theoretical folding cores

Table 1 shows the number of residues in each theoretical folding core along with the number of true positives \mathcal{T} . The comparison pictured in Figure 4 is made quantitative with the specificity α , the sensitivity γ , and the enhancement factor ϵ , as defined in Equations (1), (2) and (3), respectively. Each of the five theoretical folding cores are significant improvements upon random selections, as demonstrated by the values of ϵ , the lowest of which is 1.33. The four theoretical methods introduced in this paper all outscore FIR with respect to this enhancement measure, with F+F scoring highest for both the CypA-CsA complex and the unbound CypA, with SeS second. SeS is highly selective and sensitive in both cases, with specificity $\alpha \geq 0.77$ and sensitivity $\gamma \geq 0.74$. Unfortunately, the simple SeS method of folding core selection is clearly inappropriate for predicting changes upon ligand binding, since by definition it does not change. As discussed in Ref. (1), FIR is sensitive to the bond network and its increase in size upon ligand removal is in disagreement with the experimental results, which show a larger folding core in the presence




Figure 4: Comparison of the experimentally determined HDX folding core for CypA-CsA (blue and shown above and below the other folding cores) with the five folding cores computed using the theoretical approaches outlined in the text and labeled accordingly (from top to bottom: SeS (yellow), FIR (red), FIR_B (orange), FRO (purple), and F+F (pink)). In each case the folding core residues in the primary structure of CypA are represented as colored blocks while non-folding core residues are kept white. The residue numbers along the protein backbone are indicated horizontally, with thin black vertical lines added every ten residues for clarity.




Figure 5: Cartoon representation of a comparison of HDX folding cores with F+F for (a) the CypA-CsA complex and (b) unbound CypA. In both cases, residues are colored purple (part of both folding cores), blue (HDX folding core only), red (F+F folding core only) or green (not part of either folding core). CsA is colored black in (a). The arrow in (a) indicates residues 86-90.

Table 1: Quantitative measures of agreement between theoretical folding cores and the experimental HDX folding core for (top) CypA-CsA and (bottom) unbound CypA with $N = 165$. The bold numbers indicate the largest selectivity α , sensitivity γ and enhancement ϵ for CypA-CsA and unbound CypA. The error estimates for α and γ reflect an assumed $\pm 5\%$ variation in N_{Ex} , N_{Th} and \mathcal{T} . The two estimates for the enhancement ϵ show the variation when using the full set of residues with $N = 165$ and when excluding the 6 proline residues of CypA such that $N = 159$.

	Folding core	N_{Ex}	N_{Th}	\mathcal{T}	α	γ	ϵ_{165}	ϵ_{159}
CypA-CsA	SeS	80	75	59	0.79 ± 0.04	0.74 ± 0.04	1.62	1.56
	FIR	80	86	57	0.66 ± 0.04	0.71 ± 0.04	1.37	1.32
	FIR_B	80	72	53	0.74 ± 0.04	0.66 ± 0.04	1.52	1.46
	FRO	80	80	56	0.70 ± 0.04	0.70 ± 0.04	1.44	1.39
	F+F	80	54	44	0.81 ± 0.05	0.55 ± 0.03	1.68	1.62
unbound CypA	SeS	73	75	58	0.77 ± 0.04	0.79 ± 0.04	1.75	1.68
	FIR	73	112	66	0.59 ± 0.03	0.90 ± 0.05	1.33	1.28
	FIR_B	73	72	52	0.72 ± 0.04	0.71 ± 0.04	1.63	1.57
	FRO	73	76	54	0.71 ± 0.04	0.74 ± 0.04	1.61	1.54
	F+F	73	59	51	0.86 ± 0.05	0.70 ± 0.04	1.95	1.88

of CsA. The most accurate method for CypA-CsA and unbound CypA is F+F, with $\alpha = 0.81$ and 0.86 . This high specificity comes at the cost of lower sensitivity, meaning that it predicts with high accuracy but does not capture all of the HDX folding core. A larger proportion of the HDX folding core needs to be accurately determined by this method in order to successfully capture the impact of ligand binding on HDX. We have also estimated the variation in α , γ and ϵ in Table 1, assuming a $\pm 5\%$ error in N_{Ex} , N_{Th} and \mathcal{T} . This leads to some variation in the results for α and γ and allows us to estimate the accuracy of the numbers quoted here. In addition, due to their lack of an amide proton, the six proline residues in CypA do not appear in the ^1H - ^{15}N HSQC spectrum and so cannot be part of the HDX folding core. For this reason, we also take $N = 159$ when calculating ϵ . We find that in most cases the largest values of ϵ , α and γ remain significantly larger. This shows that our findings are indeed robust.

Effect of ligand binding on folding cores

As demonstrated in Ref. (1), for CypA there are only small changes to the HDX folding core upon ligand binding, and these are not captured using rigidity analysis alone. Indeed, only seven residues are part of the HDX folding core for the complex but not for the unbound protein. Upon ligand removal, FIR_B changes by a single residue, which is not one of these seven. As discussed in Ref. (1), the FIR changes unexpectedly upon ligand binding. These two folding cores solely derived using FIRST do not reflect the effect observed experimentally. We see a larger effect of ligand binding on FRO. Figure 6

shows how ligand binding affects the burial distance R . We calculated the average R for each residue from the final conformers of the FRODA simulations, and plotted the absolute difference, ΔR , between these values. Of




Figure 6: The difference in burial distance ΔR (black bars) between the CypA-CsA complex and the unbound protein plotted for each amino acid. The arrows and vertical dotted lines indicate the binding residues. As in Figures 2 and 1, FIR_B is also shown.

the eleven residues for which $\Delta R > 0.5$ Å, five are binding site residues and none are further than two residues along the backbone from a binding site residue. We also calculated ΔR for each residue from FRODA simulations carried out at $E_{\text{cut}} = -0.5, -1.0, -1.5$ and -3.0 kcal/mol. In each case, the residues which are part of, or close to, the binding site have the largest ΔR (data not shown). The HDX folding core changes in sections

of the protein more distant from the binding site, and so tracking the protein surface in FRODA does not capture this effect.

CONCLUSION

We have introduced and discussed rapid computational methods for adding information on surface exposure and protein dynamics into a rigidity-based folding core definition. These improve upon the FIRST folding core for predicting the HDX folding cores of both CypA-CsA and unliganded CypA. In addition, we show that they are less sensitive to the erroneous increase in folding core size upon ligand removal observed using FIRST with the standard bond network. On balance, F+F, which combines rigidity and motion, appears as the best choice for a computational determination of the folding core from a protein’s structural information alone. Nevertheless, no method achieves a perfect score of $\alpha = \gamma = 1$, neither for the uncomplexed CypA nor for the CypA-CsA complex. This insensitivity to ligand binding is disappointing but unsurprising due to the scale of the challenge — for CypA-CsA and in general.

We have sought to obtain a practical balance between the computational demands of our predictions and their utility. The requirement for large-scale motions make more computationally sophisticated methods such as MD inappropriate. Methods using simpler input data than ours, making sequence-based predictions of surface exposure and then in turn predictions of HDX folding cores, are currently unsatisfactory. We feel that although balance between sophistication of input data and method and the accuracy of the prediction is yet to be struck, it may be achieved through coarse-grained approaches. The FRODA alternative FRODAN has more flexibility of hydrogen bond and hydrophobic contact geometry built in to its bond networks. A comparison between the predictions made using FRODAN instead of FRODA would be a valuable extension to this work.

Our overall strategy is based on using the results from HDX experiments as the benchmark against which theoretical results are compared. The measures α , γ and ϵ score most highly when the match of a theoretical folding core with the HDX result is perfect. However, the HDX data may contain errors themselves and this additional source of variation did not play a role in our judgements of the success of the computational approaches. Furthermore, for the CypA-CsA complex, the protein is large relative to the ligand and the ligands effect on HDX is rather small. We expect that an application of our methods to a larger protein that exhibits a more significant conformational change upon ligand

binding will also give a clearer change in the values of α , γ and ϵ .

ACKNOWLEDGMENTS

We thank C. Blindauer for help with the experimental NMR aspects of this work. The CypA plasmid was kindly provided by G. Fischer from the Max Planck Research Unit for Enzymology of Protein Folding in Halle, Germany. We gratefully acknowledge funding from the EPSRC Life Sciences Interface programme (MOAC DTC EP/F500378/1). JWH thanks the Institute of Advanced Study in Warwick for an Early Career Fellowship; SAW is grateful to the Leverhulme Trust for an Early Career Fellowship.

References

- [1] Heal, J. W., R. A. Römer, C. A. Blindauer, and R. B. Freedman, 2014. Characterising the folding core of the cyclophilin A – cyclosporin A complex I: hydrogen exchange data and rigidity analysis. *bpj* Submitted.
- [2] Hespenheide, B. M., A. J. Rader, M. F. Thorpe, and L. A. Kuhn, 2002. Identifying Protein Folding Cores: Observing the Evolution of Rigid and Flexible Regions during Unfolding. *J. Mol. Graph. & Model.* 21:195–207.
- [3] Rader, A. J., and I. Bahar, 2004. Folding core predictions from network models of proteins. *Polymer* 45:659–668.
- [4] Rader, A. J., G. Anderson, B. Isin, H. G. Khorana, I. Bahar, and J. Klein-Seetharaman, 2004. Identification of core amino acids stabilizing rhodopsin. *Proc. Nat. Acad. Sci.* 101:7246–7252.
- [5] Woodward, C., 1993. Is the slow-exchange core the protein folding core? *Trends Biochem. Sci.* 18:359–360.
- [6] Li, R., and C. Woodward, 1999. The hydrogen exchange core and protein folding. *Prot. Sci.* 8:1571–1591.
- [7] Henzler-Wildman, K., and D. Kern, 2007. Dynamic Personalities of Proteins. *Nature* 450:964–972.
- [8] Gohlke, H., and M. F. Thorpe, 2006. A natural coarse graining for simulating large biomolecular motion. *Biophys. J.* 91:2115–2120.
- [9] Karplus, M., and G. A. Petsko, 1990. Molecular dynamics simulations in biology. *Nature* 347:631–639.
- [10] Wells, S. A., S. Menor, B. M. Hespenheide, and M. F. Thorpe, 2005. Constrained geometric simulation of diffusive motion in proteins. *Phys. Biol.* 2:S127–S136.
- [11] Clementi, C., 2008. Coarse-grained models of protein folding: toy models or predictive tools? *Curr. Opin. Struct. Biol.* 18:10–15.
- [12] Jolley, C. C., S. A. Wells, B. M. Hespenheide, M. F. Thorpe, and P. Fromme, 2006. Docking of Photosystem I Subunit C Using a Constrained Geometric Simulation. *J. Am. Chem. Soc.* 128:8803–8812.
- [13] Macchiarulo, A., R. Nuti, D. Bellochi, E. Camaioni, and R. Pellicciari, 2007. Molecular docking and spatial coarse graining simulations as tools to investigate substrate recognition, enhancer binding and conformational transitions in indoleamine-2,3-dioxygenase (IDO). *BBA-Proteins Proteom* 1774:1058–1068.

- [14] Jimenez-Roldan, J. E., R. B. Freedman, R. A. Römer, and S. A. Wells, 2012. Rapid simulation of protein motion: merging flexibility, rigidity and normal mode analyses. *Phys. Biol.* 9:016008.
- [15] Diamond, R., 1990. On the Use of Normal Modes in Thermal Parameter Refinement: theory and Application to the Bovine Pancreatic Trypsin Inhibitor. *Acta Cryst.* A46:425–435.
- [16] Hinsen, K., 1998. Analysis of domain motions by approximate normal mode calculations. *Proteins* 33:417–429.
- [17] Suhre, K., and Y.-H. Sanejouand, 2004. ElNémo: a normal mode web server for protein movement analysis and the generation of templates for molecular replacement. *Nucleic Acids Res. (Web Issue)* 32:610–614.
- [18] Petrone, P., and P. VS, 2006. Can conformational change be described by only a few normal modes? *Biophys J.* 90:p1583–1593.
- [19] Dobbins, S. E., V. I. Lesk, and M. J. E. Sternberg, 2008. Insights into protein flexibility: The relationship between normal modes and conformational change upon protein–protein docking. *P. Natl. A. Sci. USA* 105:10390–10395.
- [20] Krebs, W. G., V. Alexandrov, C. A. Wilson, N. Echols, H. Yu, and M. Gerstein, 2002. Normal mode analysis of macromolecular motions in a database framework: developing mode concentration as a useful classifying statistic. *Proteins* 48:682–695.
- [21] Alexandrov, V., U. Lehnert, N. Echols, D. Milburn, D. Engelman, and M. Gerstein, 2005. Normal modes for predicting protein motions: a comprehensive database assessment and associated Web tool. *Prot. Sci.* 14:633–643.
- [22] Sljoka, A., and D. Wilson, 2013. Probing Protein Ensemble Rigidity and Hydrogen-Deuterium exchange. pre-print.
- [23] Mikol, V., J. Kallen, G. Pflügl, and M. D. Walkinshaw, 1993. X-ray structure of a monomeric cyclophilin A-cyclosporin A crystal complex and 2.1 Å resolution. *J. Mol. Biol.* 234:1119–1130.
- [24] Folch, B., M. Rooman, and Y. Dechouck, 2008. Thermostability of salt bridges versus hydrophobic interactions in proteins probed by statistical potentials. *J. Chem. Inf. Model.* 48:119–127.
- [25] Dill, K. A., 1990. Dominant forces in protein folding. *Biochemistry* 29:7133–7155.
- [26] Thorpe, M. F., B. M. Hespeneide, Y. Yang, and L. A. Kuhn, 2000. Flexibility and Critical Hydrogen Bonds in Cytochrome c. *Pac. Symp. Biocomput.* 191–202.
- [27] Jacobs, D., A. Rader, L. Kuhn, and M. Thorpe, 2001. Protein flexibility predictions using graph theory. *Prot: Struct. Func. Gen.* 44:150–165.
- [28] Zavodszky, M. I., M. Lei, M. F. Thorpe, A. R. Day, and L. A. Kuhn, 2004. Modeling Correlated Main-Chain Motions in Proteins for Flexible Molecular Recognition. *Prot: Struct. Func. Gen.* 57:243–261.
- [29] Radestock, S., and H. Gohlke, 2008. Exploiting the Link between Protein Rigidity and Thermostability for Data-Driven Protein Engineering. *Eng. Life Sci.* 8:507–522.
- [30] Rathi, P. C., S. Radestock, and H. Gohlke, 2012. Thermostabilizing mutations preferentially occur at structural weak spots with a high mutation ratio. *J. Biotechnol.* 159:135–144.
- [31] Pflieger, C., P. C. Rathi, D. L. Klein, S. Radestock, and H. Gohlke, 2013. Constraint Network Analysis (CNA): a Python software package for efficiently linking biomacromolecular structure, flexibility, (thermo-)stability, and function. *J. Chem. Inf. Model.* 53:1007–1015.
- [32] Gohlke, H., L. A. Kuhn, and D. A. Case, 2004. Change in protein flexibility upon complex formation: analysis of ras-raf using molecular dynamics and a molecular framework approach. *Prot: Struct. Func. Bioinf.* 56:332–337.
- [33] Tirion, M. M., 1996. Large amplitude elastic motions in proteins from single-parameter atomic analysis. *Phys. Rev. Lett.* 77:1905–1908.
- [34] Wells, S., J. Jimenez-Roldan, and R. Römer, 2009. Comparative analysis of rigidity across protein families. *Phys. Biol.* 6:046005–11.
- [35] Heal, J. W., 2013. Effects of ligand binding on the rigidity and mobility of proteins: a computational and experimental approach. Ph.D. thesis, University of Warwick.